

WHAT IS CLAIMED IS:

1. A system, comprising:

5 an inter-node network; and

a plurality of nodes coupled by the inter-node network, wherein each of the plurality of nodes includes a plurality of active devices, an interface configured to send and receive coherency messages on the inter-node
10 network, and an address network coupling the plurality of active devices to the interface;

wherein an active device included in a node of the plurality of nodes is configured to initiate a transaction involving a coherency unit by sending either a first
15 type of address packet or a second type of address packet on the address network dependent on whether the active device is included in a multi-node system.

2. The system of claim 1, wherein the active device is configured to send the first
20 type of address packet if the active device is included in a multi-node system, and wherein each active device included in the node having access to or ownership of the coherency unit is configured to ignore the first type of address packet;

wherein each active device included in the node having access to or ownership of
25 the coherency unit is configured to transition an access right to or an ownership responsibility for the coherency unit in response to the second type of address packet.

3. The system of claim 1, wherein the first type of address packet is a remote write
30 back address packet, wherein the active device is configured to send the remote write

back address packet if the active device is included in a multi-node system and if the coherency unit is not mapped by any memory subsystem included in the node.

4. The system of claim 3, wherein an interface included in the node is configured to
5 send a coherency message via the inter-node network to a home node for the coherency unit in response to receiving the remote write back address packet, and wherein each active device included in the node is configured to ignore the remote write back address packet.

10 5. The system of claim 4, wherein a home interface in the home node is configured to lock the coherency unit in response to the coherency message and to responsively send an additional coherency message requesting initiation of a proxy read-to-own-modified subtransaction to the interface in the node.

15 6. The system of claim 5, wherein in response to receiving the additional coherency message, the interface in the node is configured to send a proxy read-to-own-modified address packet on the address network.

7. The system of claim 6, wherein each active device included in the node having an
20 access right to the coherency unit and not having an ownership responsibility for the coherency unit is configured to invalidate the access right in response to the proxy read-to-own modified address packet.

8. The system of claim 6, wherein the active device is configured to transition an
25 ownership responsibility for the coherency unit upon receipt of the proxy read-to-own modified address packet and to responsively send a data packet corresponding to the coherency unit to the interface.

9. The system of claim 8, wherein the active device is configured to transition an
30 access right corresponding to the coherency unit upon sending the data packet.

10. The system of claim 1, wherein the transaction is a write back transaction,
wherein the active device is configured to send the first type of address packet if the
active device is included in a multi-node system and the second type of address packet if
5 the active device is included in a single node system;

wherein if the active device sends the first type of address packet and another
active device included in the node gains ownership of the coherency unit
before an interface included in the node sends a responsive address packet,
10 the other active device is configured to provide data to the interface in
response to the responsive address packet;

wherein if the active device sends the second type of address packet and the other
active device included in the node gains ownership of the coherency unit
15 before a memory subsystem included in the node sends a different
responsive address packet, the active device is configured to send a NACK
data packet to the memory subsystem.

11. The system of claim 1, wherein the active device includes a mode register
20 configured to store a value indicating whether the active device is included in a multi-
node system.

12. The system of claim 1, wherein the first type of address packet is a remote write
stream packet, wherein the active device is configured to send the first type of address
25 packet if the active device is included in a multi-node system and to send the second type
of address packet if the active device is included in a single node system.

13. The system of claim 12, wherein an interface included in the node is configured to
respond to the first type of address packet by sending a coherency message via the inter-
30 node network to a home node for the coherency unit, and wherein active devices and

memory subsystems included in the node are configured to ignore the first type of address packet.

14. The system of claim 13, wherein in response to the coherency message, a home
5 interface included in the home node is configured to lock the coherency unit and to
responsively send an invalidating coherency message to one or more ones of the plurality
of nodes and to send a write stream coherency message to the interface in the node.

15. The system of claim 14, wherein the interface in the node is configured to send a
10 pull request data packet to the active device in response to receiving acknowledgment
coherency messages from each of the one or more ones of the plurality of nodes that
received the invalidating coherency message;

wherein in response to the pull request data packet, the active device is configured
15 to send an additional data packet containing a copy of the coherency unit
to the interface.

16. A node, comprising:

20 a plurality of devices coupled by an address network, wherein the plurality of
devices includes an active device and an interface configured to
communicate via an inter-node network coupling nodes in a multi-node
system;

25 wherein the active device is configured to initiate a transaction involving a
coherency unit by sending either a first type of address packet or a second
type of address packet on the address network dependent on whether the
active device is included in a multi-node system or a single node system.

17. The node of claim 16, wherein the active device is configured to send the first type of address packet if the active device is included in a multi-node system, and wherein each active device included in the node having access to or ownership of the coherency unit is configured to ignore the first type of address packet;

5

wherein each active device included in the node having access to or ownership of the coherency unit is configured to transition an access right to or an ownership responsibility for the coherency unit in response to the second type of address packet.

10

18. The node of claim 16, wherein the first type of address packet is a remote write back address packet, wherein the active device is configured to send the remote write back address packet if the active device is included in a multi-node system and if the coherency unit is not mapped by any memory subsystem included in the node.

15

19. The node of claim 18, wherein the interface is configured to send a coherency message via the inter-node network to a home node for the coherency unit in response to receiving the remote write back address packet, and wherein each active device included in the node is configured to ignore the remote write back address packet.

20

20. The node of claim 19, wherein in response to receiving a responsive coherency message from the home node for the coherency unit, the interface in the node is configured to send a proxy read-to-own-modified address packet on the address network.

25

21. The node of claim 20, wherein each active device included in the node having an access right to the coherency unit and not having an ownership responsibility for the coherency unit is configured to invalidate the access right in response to the proxy read-to-own modified address packet.

22. The node of claim 20, wherein the active device is configured to transition an ownership responsibility for the coherency unit upon receipt of the proxy read-to-own modified address packet and to responsively send a data packet corresponding to the coherency unit to the interface.

5

23. The node of claim 22, wherein the active device is configured to transition an access right corresponding to the coherency unit upon sending the data packet.

24. The node of claim 16, wherein the transaction is a write back transaction, wherein
10 the active device is configured to send the first type of address packet if the active device is included in a multi-node system and the second type of address packet if the active device is included in a single node system;

wherein if the active device sends the first type of address packet and another
15 active device included in the node gains ownership of the coherency unit before the interface sends a responsive address packet, the other active device is configured to provide data to the interface in response to the responsive address packet;

20 wherein if the active device sends the second type of address packet and the other active device included in the node gains ownership of the coherency unit before a memory subsystem included in the node sends a different responsive address packet, the active device is configured to send a NACK data packet to the memory subsystem.

25

25. The node of claim 16, wherein the active device includes a mode register configured to store a value indicating whether the active device is included in a multi-node system.

26. The node of claim 16, wherein the first type of address packet is a remote write stream packet, wherein the active device is configured to send the first type of address packet if the active device is included in a multi-node system and to send the second type of address packet if the active device is included in a single node system.

5

27. The node of claim 26, wherein the interface is configured to respond to the first type of address packet by sending a coherency message via the inter-node network to a home node for the coherency unit, and wherein active devices and memory subsystems included in the node are configured to ignore the first type of address packet.

10

28. A method of operating a multi-node system comprising a plurality of nodes coupled by an inter-node network, wherein each of the plurality of nodes includes a plurality of active devices, an interface configured to send and receive coherency messages on the inter-node network, and an address network coupling the plurality of active devices to the interface, the method comprising:

15

an active device included in a node of the plurality of nodes detecting whether the active device is included in a multi-node system; and

20

an active device included in a node of the plurality of nodes initiating a transaction involving a coherency unit by sending either a first type of address packet or a second type of address packet on an address network included in the node dependent on said detecting.

25

29. The method of claim 28, wherein the active device sends the first type of address packet if the active device is included in a multi-node system, the method further comprising:

30

each active device included in the node having access to or ownership of the coherency unit ignoring the first type of address packet; and

5 each active device included in the node having access to or ownership of
the coherency unit transitioning an access right to or an ownership
responsibility for the coherency unit in response to the second type of
address packet.

30. The method of claim 28, wherein the first type of address packet is a remote write
back address packet, the method further comprising:

10 the active device sending the remote write back address packet if the active device
is included in a multi-node system and if the coherency unit is not mapped
by any memory subsystem included in the node.

31. The method of claim 30, further comprising:

15 an interface included in the node sending a coherency message via the inter-node
network to a home node for the coherency unit in response to receiving the
remote write back address packet; and

20 each active device included in the node ignoring the remote write back address
packet.

32. The method of claim 31, further comprising a home interface in the home node
locking the coherency unit in response to the coherency message and responsively
25 sending an additional coherency message requesting initiation of a proxy read-to-own-
modified subtransaction to the interface in the node.

33. The method of claim 32, further comprising the interface in the node sending a
proxy read-to-own-modified address packet on the address network in response to
30 receiving the additional coherency message,.

34. The method of claim 33, further comprising each active device included in the node having an access right to the coherency unit and not having an ownership responsibility for the coherency unit invalidating the access right in response to the proxy
5 read-to-own modified address packet.

35. The method of claim 33, further comprising the active device transitioning an ownership responsibility for the coherency unit upon receipt of the proxy read-to-own modified address packet and responsively sending a data packet corresponding to the
10 coherency unit to the interface.

36. The method of claim 37, further comprising the active device transitioning an access right corresponding to the coherency unit upon sending the data packet.

15 37. The method of claim 28, wherein the transaction is a write back transaction, the method further comprising:

the active device sending the first type of address packet if the active device is included in a multi-node system and the second type of address packet if
20 the active device is included in a single node system;

if the active device sends the first type of address packet and another active device included in the node gains ownership of the coherency unit before an interface included in the node sends a responsive address packet, the other
25 active device providing data to the interface in response to the responsive address packet;

if the active device sends the second type of address packet and the other active device included in the node gains ownership of the coherency unit before a
30 memory subsystem included in the node sends a different responsive

address packet, the active device sending a NACK data packet to the memory subsystem.

38. The method of claim 28, further comprising a mode register included in the active device storing a value indicating whether the active device is included in a multi-node system.

39. The method of claim 28, wherein the first type of address packet is a remote write stream packet, the method further comprising:

10

the active device sending the first type of address packet if the active device is included in a multi-node system and sending the second type of address packet if the active device is included in a single node system.

15 40. The method of claim 39, further comprising:

an interface included in the node responding to the first type of address packet by sending a coherency message via the inter-node network to a home node for the coherency unit, wherein active devices and memory subsystems included in the node ignore the first type of address packet.

20

41. The method of claim 40, further comprising:

in response to the coherency message, a home interface included in the home node locking the coherency unit and responsively sending an invalidating coherency message to one or more ones of the plurality of nodes and sending a write stream coherency message to the interface in the node.

25

42. The method of claim 41, further comprising:

30

the interface in the node sending a pull request data packet to the active device in response to receiving acknowledgment coherency messages from each of the one or more ones of the plurality of nodes that received the invalidating coherency message;

5

in response to the pull request data packet, the active device sending an additional data packet containing a copy of the coherency unit to the interface.

10